Multi-Level Parallelism: A New Computing Technique

James R. Taft

High-performance computing platforms are evolving toward systems with ever larger numbers of processors. These systems utilize standard off-the-shelf microprocessors at the heart of their design. Virtually all hardware vendors have adopted this design approach because it dramatically reduces their costs. Unfortunately, this usually forces researchers to embark on large code-conversion efforts (often consuming work-years) to take advantage of the new parallel systems. Codes used in heavy production environments were often deemed impossible to convert.

Several events have occurred that are changing that attitude. First, the vendors are now building systems that share many of the hardware features of the Cray vector systems. In particular, the new designs are moving toward true shared memory architectures, and instruction sets compatible with vector programming constructs. These new hardware attributes have made it possible to approach parallelism in an entirely new way, a way that is more intuitive and simpler to implement.

Ames Research Center has extensively investigated these new approaches to parallelism using the Silicon Graphics, Inc. Origin 2000 as a test bed. The end result is a new technique for parallel computing: multi-level parallelism (MLP). This project demonstrated that new parallel systems, such as the Origin 2000, could be used in a heavy computational fluid dynamics (CFD) production environment. Furthermore, the CFD codes, running on the Cray C90 systems, could be converted to the parallel systems and still maintain their high-performance levels on the vector machines.

The production CFD code OVERFLOW was chosen as the test bed for the MLP effort. OVERFLOW is composed of approximately 100,000 lines of FORTRAN. It was chosen because it represented one of the toughest tests for the resiliency of the MLP programming approach. To further stress-test the method, it was decided to execute only the largest and most complex problems. To this end a 33-million-point problem was selected. This problem was the largest that had ever been solved at Ames. It requires hundreds of C90 hours to converge.

Doing the zones in parallel is not new; the Message Passing Interface version of OVERFLOW already attempts this process. The unique feature of the MLP approach is that it does so with no message passing and only a few hundred lines of code changes. This results in a code that is simple to maintain, that continues to execute well on C90 systems, and that now executes well on parallel systems at very high sustained levels of performance.

The MLP version of OVERFLOW was completed in about 3 workweeks. The initial performance of the code for the 33-million-point problem was 75 percent that of the performance of the dedicated 16 CPU C90 system when executed on a 64 CPU Origin system. This amounts to 3.7 GFLOPS (10^9 floating point operations per second). This relative performance coupled with the dramatically smaller price for the Origin 2000 resulted in a price/performance advantage of over 13.5. Further tests have demonstrated a sustained performance of over 6.3 GFLOPS for this problem on a 128 CPU Origin, with performance scaling still virtually linear at 128 processors. This performance level is approximately 40% better than that of the dedicated C90.

One of the most important features of the MLP technique is that the user can simply comment-out a few lines of code that call the MLP routines, and the code reverts back to the standard vector-friendly C90 version. This minimal change from the C90 version means that new releases of the OVERFLOW code can be converted rapidly to MLP. This was recently demonstrated with the introduction of OVERFLOW 1.8 which was converted to MLP form in less than 1 day.

The MLP technique has general applicability to many CFD codes. MLP offers researchers the possibility of rapidly converting their Cray vector codes to the new machines, and of insuring high performance on each. Even more importantly, NASA researchers are now able to use the new cost-effective parallel platforms to solve problems, the solutions to which were previously unaffordable.

Point of Contact: J. Taft
(650) 604-0704
jtaft@nas.nasa.gov